Google

Web Images Groups News Froogle Local **more »**

logical processor idle and logical processor bu Search Advanced Search Preferences

O Search the Web ◉ Search English pages

The **"AND"** operator is unnecessary – we include all search terms by default. [details]

Web Results **1** - **10** of about **196,000 English** pages for **logical processor idle and logical processor busy**. (**0.44** se

CPU FAQs
Processor busy is measured using an **Idle** thread mechanism. ... Since the assignment of **logical processor** numbers to physical **processor** cores is a BIOS ...
www.demandtech.com/FAQsCPU.htm - 43k - Cached - Similar pages

Intel® Software Network - Impact of Load Imbalance on **Processors** ...
However, each physical **processor** contains multiple **logical processors** and each of those ... In particular, applications should not use **busy**-wait loops for ...
www.intel.com/cd/ids/developer/ asmo-na/eng/20477.htm?prn=Y - 22k - Cached - Similar pages

Intel® Pentium® 4 **Processor** - Integration Overview for Systems ...
Comparison of a Pentium® 4 **Processor** Supporting Hyper-Threading ... that uses two separate physical **processors** (see Figure 1), the **logical processors** in a ...
www.intel.com/support/ **processors**/pentium4/pentium4_ht.htm - 59k - Aug 10, 2005 - Cached - Similar pages

Win32_PerfRawData_PerfOS_Processor class [WMI]
... is the part of the computer that performs arithmetic and **logical** computations, ... Percentage of non-**idle processor** time spent in privileged mode. ...
msdn.microsoft.com/library/en-us/ wmisdk/wmi/win32_perfrawdata_perfos_**processor**.asp - 38k - Cached - Similar pages

Intel Dual Core: Multi-Tasking Benchmarking
Keep in mind that Hyperthreading is not an entire **logical processor**, ... and a **logical processor**, and will try to schedule threads to a **idle** physical ...
forums.legitreviews.com/about1669-20.html - 66k - Cached - Similar pages

LostCircuits, CPU Guide
Everybody knows that not every workday is 100% **busy**, rather, there are phases of ... The CPU usage of both **logical processors** adds up to 100%, however, ...
www.lostcircuits.com/cpu/p4_306/2.shtml - 21k - Cached - Similar pages

Hyper-Threading Linux @ LINUXWORLD MAGAZINE
... sit **idle**, while the **processor** reports itself as **busy** to the operating system. ... In an SMT system, because the **logical processors** share cache, ...
linux.sys-con.com/read/33885.htm - 83k - Aug 10, 2005 - Cached - Similar pages

Hyper-threading - Enpsychlopedia
... that are both Hyper-Threaded (for a total of four **logical processors**). ... one CPU would be extremely **busy** while the other CPU would be completely **idle**, ...
psychcentral.com/psypsych/Hyperthreading - Similar pages

sar Command
If no disk I/O is in progress and the CPU is not **busy**, the **idle** category gets the ... On a two-**logical processor** system, this produces output similar to the ...
publib.boulder.ibm.com/infocenter/ pseries/topic/com.ibm.aix.doc/cmds/aixcmds5/sar.htm - 38k - Cached - Similar pages

Computing.Net - HT & HTT
... CPU would be extremely **busy** while the other CPU would be completely **idle**, ... Thats really cool...it also seems more **logical** that the **processor** would ...
www.computing.net/cpus/wwwboard/forum/11742.html - 29k - Cached - Similar pages

PORTAL USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

performance workloads in a hardware multi threaded environm

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used
performance workloads in a hardware multi threaded environment

Found **58,149** of **119,996**

Sort results by: publication date

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 101 - 120 of 200 Result page: previous 1 2 3 4 5 **6** 7 8 9 10 next

Best 200 shown Relevance scale

101 Hardware assisted unstructured volume rendering: Parallelizing a high accuracy hardware-assisted volume renderer for meshes with arbitrary polyhedra

Janine Bennett, Richard Cook, Nelson Max, Deborah May, Peter Williams

October 2001 **Proceedings of the IEEE 2001 symposium on parallel and large-data visualization and graphics**

Full text available: pdf(188.70 KB) Additional Information: full citation, abstract, references, index terms

This paper discusses our efforts to improve the performance of the high-accuracy (HIAC) volume rendering system, based on cell projection, which is used to display unstructured, scientific data sets for analysis. The parallelization of HIAC, using the pthreads and MPI API's, resulted in significant speedup, but interactive frame rates are not yet attainable for very large data sets.

102 Computing curricula 2001

September 2001 **Journal on Educational Resources in Computing (JERIC)**

Full text available: pdf(613.63 KB) html(2.78 KB) Additional Information: full citation, references, citings, index terms

103 A user-programmable vertex engine

Erik Lindholm, Mark J. Kligard, Henry Moreton

August 2001 **Proceedings of the 28th annual conference on Computer graphics and interactive techniques**

Full text available: pdf(12.03 MB) Additional Information: full citation, abstract, references, citings, index terms

In this paper we describe the design, programming interface, and implementation of a very efficient user-programmable vertex engine. The vertex engine of NVIDIA's GeForce3 GPU evolved from a highly tuned fixed-function pipeline requiring considerable knowledge to program. Programs operate only on a stream of independent vertices traversing the pipe. Embedded in the broader fixed function pipeline, our approach preserves parallelism sacrificed by previous approaches. The programmer is presente ...

Keywords: graphics hardware, graphics systems

104 Implicit coscheduling: coordinated scheduling with implicit information in distributed systems

Andrea Carol Arpaci-Dusseau

August 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 3

Full text available: pdf(1.83 MB) Additional Information: full citation, abstract, references, citings, index terms

In modern distributed systems, coordinated time-sharing is required for communicating processes to leverage the performance of switch-based networks and low-overhead protocols. Coordinated time-sharing has traditionally been achieved with gang scheduling or explicit coscheduling, implementations of which often suffer from many deficiencies: multiple points of failure, high context-switch overheads, and poor interaction with client-server, interactive, and I/O -intensive workloads. I ...

Keywords: clusters, coscheduling, gang scheduling, networks of workstations, proportional-share scheduling, two-phase waiting

105 Towards a first vertical prototyping of an extremely fine-grained parallel programming approach

Dorit Naishlos, Joseph Nuzman, Chau-Wen Tseng, Uzi Vishkin

July 2001 **Proceedings of the thirteenth annual ACM symposium on Parallel algorithms and architectures**

Full text available: pdf(341.17 KB) Additional Information: full citation, abstract, references, citings, index terms

Explicit-multithreading (XMT) is a parallel programming approach for exploiting on-chip parallelism. XMT introduces a computational framework with 1) a simple programming style that relies on fine-grained PRAM-style algorithms; 2) hardware support for low-overhead parallel threads, scalable load balancing, and efficient synchronization. The missing link between the algorithmic-programming level and the architecture level is provided by the first prototype XMT compiler. This paper also takes t ...

Keywords: compilers, parallel programming, processor architecture

106 Cache performance for multimedia applications

Nathan T. Slingerland, Alan Jay Smith

June 2001 **Proceedings of the 15th international conference on Supercomputing**

Full text available: pdf(642.63 KB) Additional Information: full citation, abstract, references, citings, index terms

The caching behavior of multimedia applications has been described as having high instruction reference locality within small loops, very large working sets, and poor data cache performance due to non-locality of data references. Despite this, there is no published research deriving or measuring these qualities. Utilizing the previously developed Berkeley Multimedia Workload, we present the results of execution driven cache simulations with the goal of aiding future media processing architect ...

Keywords: CPU caches, cache, mulitmedia, simulation, trace driven simulation

107 α-coral: a multigrain, multithreaded processor architecture

Mark N. Yankelevsky, Constantine D. Polychronopoulos

June 2001 **Proceedings of the 15th international conference on Supercomputing**

Full text available: pdf(196.56 KB) Additional Information: full citation, abstract, references, index terms

Recently popularized hardware multithreading (HMT) architectures, such as SMT, Multiscalar and Terra do not provide flexible and efficient methods of thread management and synchronization in hardware. The &agr;-Coral architecture is a tool for investigation of a more dynamic approach to thread management. Unlike other architectures, there are no strict requirements on timing and size of threads, and no static partitioning of resources. &agr;-Coral provides for simultaneous multiprogramming an ...

Keywords: multithreaded, parallelizing compiler, processor archietecture

108 Characterizing the memory behavior of Java workloads: a structured view and opportunities for optimizations

Yefim Shuf, Mauricio J. Serrano, Manish Gupta, Jaswinder Pal Singh

June 2001 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2001 ACM SIGMETRICS international conference on Measurement and modeling of computer systems**, Volume 29 Issue 1

Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, citings

This paper studies the memory behavior of important Java workloads used in benchmarking Java Virtual Machines (JVMs), based on instrumentation of both application and library code in a state-of-the-art JVM, and provides structured information about these workloads to help guide systems' design. We begin by characterizing the inherent memory behavior of the benchmarks, such as information on the breakup of heap accesses among different categories and on the hotness of references to fields and met ...

109 Power and energy reduction via pipeline balancing

R. Iris Bahar, Srilatha Manne

May 2001 **ACM SIGARCH Computer Architecture News , Proceedings of the 28th annual international symposium on Computer architecture**, Volume 29 Issue 2

Full text available: pdf(1.06 MB) Additional Information: full citation, abstract, references, citings, index terms

Minimizing power dissipation is an important design requirement for both portable and non-portable systems. In this work, we propose an architectural solution to the power problem that retains performance while reducing power. The technique, known as Pipeline Balancing (PLB), *dynamically tunes the resources of a general purpose processor to the needs of the program by monitoring performance within each program. We analyze metrics for triggering* PLB, *and detail instruction que ...*

110 Speculative precomputation: long-range prefetching of delinquent loads

Jamison D. Collins, Hong Wang, Dean M. Tullsen, Christopher Hughes, Yong-Fong Lee, Dan Lavery, John P. Shen

May 2001 **ACM SIGARCH Computer Architecture News , Proceedings of the 28th annual international symposium on Computer architecture**, Volume 29 Issue 2

Full text available: pdf(995.50 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper explores Speculative Precomputation, a technique that uses idle thread context in a multithreaded architecture to improve performance of single-threaded applications. It attacks program stalls from data cache misses by pre-computing future memory accesses in available thread contexts, and prefetching these data. This technique is evaluated by simulating the performance of a research processor based on the Itanium™ ISA supporting Simultaneous Multithreading. Two primary for ...

111 I/O reference behavior of production database workloads and the TPC benchmarks—an analysis at the logical level

Windsor W. Hsu, Alan Jay Smith, Honesty C. Young

March 2001 **ACM Transactions on Database Systems (TODS)**, Volume 26 Issue 1

Full text available: pdf(5.42 MB) Additional Information: full citation, abstract, references, citings, index terms

As improvements in processor performance continue to far outpace improvements in storage performance, I/O is increasingly the bottleneck in computer systems, especially in large database systems that manage huge amoungs of data. The key to achieving good I/O performance is to thoroughly understand its characteristics. In this article we present a comprehensive analysis of the logical I/O reference behavior of the peak

productiondatabase workloads from ten of the world's largest corporatio ...

Keywords: I/O, TPC benchmarks, caching, locality, prefetching, production database workloads, reference behavior, sequentiality, workload characterization

112 Thread-level parallelism and interactive performance of desktop applications

Kristián Flautner, Rich Uhlig, Steve Reinhardt, Trevor Mudge

November 2000 **Proceedings of the ninth international conference on Architectural support for programming languages and operating systems**, Volume 28 , 34 Issue 5 , 5

Full text available: pdf(234.58 KB) Additional Information: full citation, abstract, references, citings, index terms

Multiprocessing is already prevalent in servers where multiple clients present an obvious source of thread-level parallelism. However, the case for multiprocessing is less clear for desktop applications. Nevertheless, architects are designing processors that count on the availability of thread-level parallelism. Unlike server workloads, the primary requirement of interactive applications is to respond to user events under human perception bounds rather than to maximize end-to-end throughput. In ...

113 An analysis of operating system behavior on a simultaneous multithreaded architecture

Joshua A. Redstone, Susan J. Eggers, Henry M. Levy

November 2000 **Proceedings of the ninth international conference on Architectural support for programming languages and operating systems**, Volume 28 , 34 Issue 5 , 5

Full text available: pdf(227.80 KB) Additional Information: full citation, abstract, references, citings, index terms, review

This paper presents the first analysis of operating system execution on a simultaneous multithreaded (SMT) processor. While SMT has been studied extensively over the past 6 years, previous research has focused entirely on user-mode execution. However, many of the applications most amenable to multithreading technologies spend a significant fraction of their time in kernel code. A full understanding of the behavior of such workloads therefore requires execution and measurement of the operating sy ...

114 Thread-level parallelism and interactive performance of desktop applications

Krisztián Flautner, Rich Uhlig, Steve Reinhardt, Trevor Mudge

November 2000 **ACM SIGPLAN Notices**, Volume 35 Issue 11

Full text available: pdf(2.94 MB) Additional Information: full citation, abstract, references, index terms

Multiprocessing is already prevalent in servers where multiple clients present an obvious source of thread-level parallelism. However, the case for multiprocessing is less clear for desktop applications. Nevertheless, architects are designing processors that count on the availability of thread-level parallelism. Unlike server workloads, the primary requirement of interactive applications is to respond to user events under human perception bounds rather than to maximize end-to-end throughput. In ...

115 An analysis of operating system behavior on a simultaneous multithreaded architecture

Joshua A. Redstone, Susan J. Eggers, Henry M. Levy

November 2000 **ACM SIGPLAN Notices**, Volume 35 Issue 11

Full text available: pdf(1.56 MB) Additional Information: full citation, abstract, references, index terms

This paper presents the first analysis of operating system execution on a simultaneous multithreaded (SMT) processor. While SMT has been studied extensively over the past 6 years, previous research has focused entirely on user-mode execution. However, many of the applications most amenable to multithreading technologies spend a significant fraction of their time in kernel code. A full understanding of the behavior of such workloads

therefore requires execution and measurement of the operating sy ...

116 Concurrent garbage collection using hardware-assisted profiling
Timothy H. Heil, James E. Smith
October 2000 **ACM SIGPLAN Notices , Proceedings of the 2nd international symposium on Memory management**, Volume 36 Issue 1
Full text available: pdf(1.74 MB) Additional Information: full citation, abstract, citings, index terms

In the presence of on-chip multithreading, a Virtual Machine (VM) implementation can readily take advantage of *service threads* for enhancing performance by performing tasks such as profile collection and analysis, dynamic optimization, and garbage collection concurrently with program execution. In this context, a hardware-assisted profiling mechanism is proposed. The *Relational Profiling Architecture* (RPA) is designed from the top down. RPA is based on a relational model similar ...

117 Process migration
September 2000 **ACM Computing Surveys (CSUR)**, Volume 32 Issue 3
Full text available: pdf(1.24 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Process migration is the act of transferring a process between two machines. It enables dynamic load distribution, fault resilience, eased system administration, and data access locality. Despite these goals and ongoing research efforts, migration has not achieved widespread use. With the increasing deployment of distributed systems in general, and distributed operating systems in particular, process migration is again receiving more attention in both research and product development. As hi ...

Keywords: distributed operating systems, distributed systems, load distribution, process migration

118 Analytic model of Web servers in distributed environments
Paul Reeser, Rema Hariharan
September 2000 **Proceedings of the 2nd international workshop on Software and performance WOSP '00**
Full text available: pdf(121.84 KB) Additional Information: full citation, references, citings, index terms

Keywords: HTTP, Java, OO, Web, distributed, performance, script, servlet

119 End to End Performance Modeling of Web Sewer Architectures
R. Hadharan, W. K. Ehrlich, D. Cura, P. K. Reeser
September 2000 **ACM SIGMETRICS Performance Evaluation Review**, Volume 28 Issue 2
Full text available: pdf(1.16 MB) Additional Information: full citation, abstract

Web server performance in a distributed Object-Oriented (OO) environment is a complex interplay between a variety of factors (e.g., hardware platform, threading model, object scope model, server operating system, network bandwidth, disk file size, caching). In this paper, we present a model-based approach to Web Server performance evaluation in terms of an end-to-end queueing model implemented in a simulation tool. We have applied this model to Active Server Page (ASP) and Common Object Model (C ...

120 Cellular disco: resource management using virtual clusters on shared-memory multiprocessors
Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum
August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3
Additional Information:

Full text available: pdf(287.05 KB) full citation, abstract, references, citings, index terms, review

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that laverages the existing operating system technology. In this paper we present a ...

Keywords: fault containment, resource managment, scalable multiprocessors, virtual machines

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

PORTAL USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
performance workloads in a hardware multi threaded environm
THE ACM DIGITAL LIBRARY
Feedback Report a problem Satisfaction survey

Terms used
performance workloads in a hardware multi threaded environment

Found **58,149** of **119,996**

Sort results by: publication date
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 121 - 140 of 200 Result page: previous 1 2 3 4 5 6 **7** 8 9 10 next
Best 200 shown Relevance scale

121 High performance adaptive middleware for CORBA-based systems
E-Kai Shen, Shikharesh Majumdar, Istabrak Abdul-Fatah
July 2000 **Proceedings of the nineteenth annual ACM symposium on Principles of distributed computing**

Full text available: pdf(1.48 MB) Additional Information: full citation, abstract, references, citings, index terms

Middleware provides inter-operability and transparent location of servers in a heterogeneous distributed environment. A careful design of the middleware software is required however for achieving high performance. This research proposes an adaptive middleware architecture for CORBA-based systems. The adaptive middleware agent that maps an object name to the object reference has two modes of operations. In the handle-driven mode it returns a reference for the requested object to the ...

Keywords: CORBA performance, adaptive middleware architectures, distributed system performance, high performance middleware, middleware performance

122 Pomegranate: a fully scalable graphics architecture
Matthew Eldridge, Homan Igehy, Pat Hanrahan
July 2000 **Proceedings of the 27th annual conference on Computer graphics and interactive techniques**

Full text available: pdf(508.39 KB) Additional Information: full citation, abstract, references, citings, index terms

Pomegranate is a parallel hardware architecture for polygon rendering that provides scalable input bandwidth, triangle rate, pixel rate, texture memory and display bandwidth while maintaining an immediate-mode interface. The basic unit of scalability is a single graphics pipeline, and up to 64 such units may be combined. Pomegranate's scalability is achieved with a novel "sort-everywhere" architecture that distributes work in a balanced fashion at every stage of the pipeline, ke ...

Keywords: graphics hardware, parallel computing

123 Improving interactive performance using TIPME
Yasuhiro Endo, Margo Seltzer
June 2000 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2000 ACM SIGMETRICS international conference on Measurement and modeling of computer systems**, Volume 28 Issue 1

Additional Information:

Full text available: pdf(1.05 MB) full citation, abstract, references, citings, index terms

On the vast majority of today's computers, the dominant form of computation is GUI-based user interaction. In such an environment, the user's perception is the final arbiter of performance. Human-factors research shows that a user's perception of performance is affected by unexpectedly long delays. However, most performance-tuning techniques currently rely on throughput-sensitive benchmarks. While these techniques improve the average performance of the system, they do littl ...

Keywords: interactive performance, monitoring

124 A case for user-level dynamic page migration
Dimitrios S. Nikolopoulos, Theodore S. Papatheodorou, Constantine D. Polychronopoulos, Jesús Labarta, Eduard Ayguadé
May 2000 **Proceedings of the 14th international conference on Supercomputing**

Full text available: pdf(1.33 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents user-level dynamic page migration, a runtime technique which transparently enables parallel programs to tune their memory performance on distributed shared memory multiprocessors, with feedback obtained from dynamic monitoring of memory activity. Our technique exploits the iterative nature of parallel programs and information available to the program both at compile time and at runtime in order to improve the accuracy and the timeliness of page migration ...

125 Using complete system simulation to characterize SPECjvm98 benchmarks
Tao Li, Lizy Kurian John, Vijaykrishnan Narayanan, Anand Sivasubramaniam, Jyotsna Sabarinathan, Anupama Murthy
May 2000 **Proceedings of the 14th international conference on Supercomputing**

Full text available: pdf(1.66 MB) Additional Information: full citation, abstract, references, citings, index terms

Complete system simulation to understand the influence of architecture and operating systems on application execution has been identified to be crucial for systems design. While there have been previous attempts at understanding the architectural impact of Java programs, there has been no prior work investigating the operating system (kernel) activity during their executions. This problem is particularly interesting in the context of Java since it is not only the application that can invoke ...

126 A simulation-based study of scheduling mechanisms for a dynamic cluster environment
Yanyong Zhang, Anand Sivasubramaniam, Jose Moreira, Hubertus Franke
May 2000 **Proceedings of the 14th international conference on Supercomputing**

Full text available: pdf(1.08 MB) Additional Information: full citation, abstract, references, citings, index terms

Scheduling of processes onto processors of a parallel machine has always been an important and challenging area of research. The issue becomes even more crucial and difficult as we gradually progress to the use of off-the-shelf workstations, operating systems, and high bandwidth networks to build cost-effective clusters for demanding applications. Clusters are gaining acceptance not just in scientific applications that need supercomputing power, but also in domains such as databases, web se ...

Keywords: clusters, coscheduling, dynamic coscheduling, parallel scheduling, simulation

127 Characterizing processor architectures for programmable network interfaces
Patrick Crowley, Marc E. Fluczynski, Jean-Loup Baer, Brian N. Bershad

May 2000 **Proceedings of the 14th international conference on Supercomputing**

Full text available: pdf(984.97 KB) Additional Information: full citation, abstract, references, citings, index terms

The rapid advancements of networking technology have boosted potential bandwidth to the point that the cabling is no longer the bottleneck. Rather, the bottlenecks lie at the crossing points, the nodes of the network, where data traffic is intercepted or forwarded. As a result, there has been tremendous interest in speeding those nodes, making the equipment run faster by means of specialized chips to handle data trafficking. The Network Processor is the blanket name thrown ...

128 Smart Memories: a modular reconfigurable architecture

Ken Mai, Tim Paaske, Nuwan Jayasena, Ron Ho, William J. Dally, Mark Horowitz

May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture**, Volume 28 Issue 2

Full text available: pdf(80.16 KB) Additional Information: full citation, abstract, references, citings, index terms

Trends in VLSI technology scaling demand that future computing devices be narrowly focused to achieve high performance and high efficiency, yet also target the high volumes and low costs of widely applicable general purpose designs. To address these conflicting requirements, we propose a modular reconfigurable architecture called Smart Memories, targeted at computing needs in the 0.1&mgr; technology generation. A Smart Memories chip is made up of many processing tiles, each containing local ...

129 Piranha: a scalable architecture based on single-chip multiprocessing

Luiz André Barroso, Kourosh Gharachorloo, Robert McNamara, Andreas Nowatzyk, Shaz Qadeer, Barton Sano, Scott Smith, Robert Stets, Ben Verghese

May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture**, Volume 28 Issue 2

Full text available: pdf(191.10 KB) Additional Information: full citation, abstract, references, citings, index terms

The microprocessor industry is currently struggling with higher development costs and longer design times that arise from exceedingly complex processors that are pushing the limits of instruction-level parallelism. Meanwhile, such designs are especially ill suited for important commercial applications, such as on-line transaction processing (OLTP), which suffer from large memory stall times and exhibit little instruction-level parallelism. Given that commercial applications constitute by fa ...

130 Session summaries from the 17th symposium on operating systems principle (SOSP'99)

Jay Lepreau, Eric Eide

April 2000 **ACM SIGOPS Operating Systems Review**, Volume 34 Issue 2

Full text available: pdf(3.15 MB) Additional Information: full citation, index terms

131 A decision support system for tuning Web servers in distributed object oriented network architectures

R. D. van der Mei, W. K. Ehrlich, P. K. Reeser, J. P. Francisco

March 2000 **ACM SIGMETRICS Performance Evaluation Review**, Volume 27 Issue 4

Full text available: pdf(648.79 KB) Additional Information: full citation, abstract, citings, index terms

Web technologies are currently being employed to provide end user interfaces in diverse computing environments. The core element of these Web solutions is a Web server that is based on the Hypertext Transfer Protocol (HTTP) running over TCP/IP. Web servers are required to respond to millions of transaction requests per day at an "acceptable" Quality of Service (QoS) level with respect to the end-to-end response time and the server throughput. In many applications, the server performs significant ...

Keywords: Decision Support System, HTTP, Web server, World Wide Web, architecture, computing, configuration tuning, distributed, httpd, object-oriented, performance

132 How to solve the current memory access and data transfer bottlenecks: at the processor architecture or at the compiler level

Francky Catthoor, Nikil D. Dutt, Christoforos E. Kozyrakis

January 2000 **Proceedings of the conference on Design, automation and test in Europe**

Full text available: pdf(92.68 KB) Publisher Site

Additional Information: full citation, references, citings, index terms

133 EMERALDS: a small-memory real-time microkernel

Khawar M. Zuberi, Padmanabhan Pillai, Kang G. Shin

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles**, Volume 33 Issue 5

Full text available: pdf(1.59 MB)

Additional Information: full citation, abstract, references, citings, index terms

EMERALDS (Extensible Microkernel for Embedded, ReAL-time, Distributed Systems) is a real-time microkernel designed for small-memory embedded applications. These applications must run on slow (15-25MHz) processors with just 32-128 kbytes of memory, either to keep production costs down in mass-produced systems or to keep weight and power consumption low. To be feasible for such applications, the OS must not only be small in size (less than 20 kbytes), but also have low-overhead kernel services. Un ...

134 Cellular Disco: resource management using virtual clusters on shared-memory multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles**, Volume 33 Issue 5

Full text available: pdf(1.93 MB)

Additional Information: full citation, abstract, references, citings, index terms

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology.In this paper we present a syste ...

135 The use of multithreading for exception handling

Craig B. Zilles, Joel S. Emer, Gurindar S. Sohi

November 1999 **Proceedings of the 32nd annual ACM/IEEE international symposium on Microarchitecture**

Full text available: pdf(1.49 MB) Publisher Site

Additional Information: full citation, abstract, references, citings, index terms

Common hardware exceptions, when implemented by trapping, unnecessarily serialize program execution in dynamically scheduled superscalar processors. To avoid the consequences of trapping the main program thread, multithreaded CPUs can exploit control and data independence by executing the exception handler in a separate hardware context. The main thread doesn't squash instructions after the excepting instruction, conserving fetch bandwidth and allowing execution of instructions inde ...

136 Cluster I/O with River: making the fast case common
Remzi H. Arpaci-Dusseau, Eric Anderson, Noah Treuhaft, David E. Culler, Joseph M. Hellerstein, David Patterson, Kathy Yelick
May 1999 **Proceedings of the sixth workshop on I/O in parallel and distributed systems**
Full text available: pdf(1.20 MB) Additional Information: full citation, references, citings, index terms

137 Design challenges of virtual networks: fast, general-purpose communication
Alan M. Mainwaring, David E. Culler
May 1999 **ACM SIGPLAN Notices , Proceedings of the seventh ACM SIGPLAN symposium on Principles and practice of parallel programming**, Volume 34 Issue 8
Full text available: pdf(1.57 MB) Additional Information: full citation, abstract, references, citings, index terms

Virtual networks provide applications with the illusion of having their own dedicated, high-performance networks, although network interfaces posses limited, shared resources. We present the design of a large-scale virtual network system and examine the integration of communication programming interface, system resource management, and network interface operation. Our implementation on a cluster of 100 workstations quantifies the impact of virtualization on small message latencies and throughput ...

Keywords: application programming interfaces, direct network access, high-performance clusters, protocol architecture and implementation, system resource management, virtual networks

138 Scheduling issues in high-performance computing
L. W. Dowdy, E. Rosti, G. Serazzi, E. Smirni
March 1999 **ACM SIGMETRICS Performance Evaluation Review**, Volume 26 Issue 4
Full text available: pdf(974.37 KB) Additional Information: full citation, abstract, citings, index terms

In this paper we consider the problem of scheduling computational resources across a range of high-performance systems, from tightly coupled parallel systems to loosely coupled ones like networks of workstations and geographically dispersed meta-computing environments. We review the role of architecture issues in the choice of scheduling discipline and we present a selected set of policies that address different aspects of the scheduling problem. This discussion serves as the motivation for addr ...

139 Communication overlap in multi-tier parallel algorithms
Scott B. Baden, Stephen J. Fink
November 1998 **Proceedings of the 1998 ACM/IEEE conference on Supercomputing (CDROM)**
Full text available: pdf(278.73 KB) Additional Information: full citation, abstract, references, citings

Hierarchically organized multicomputers such as SMP clusters offer new opportunities and new challenges for high-performance computation, but realizing their full potential remains a formidable task. We present a hierarchical model of communication targeted to block-structured, bulk-synchronous applications running on dedicated clusters of symmetric multiprocessors. Our model supports node-level rather processor-level communication as the fundamental operation, and is optimized for aggregate pat ...

140 A wideband approach to integrating performance prediction into a software design environment
Murray Woodside, Curtis Hrischuk, Bran Selic, Stefan Bayarov
October 1998 **Proceedings of the 1st international workshop on Software and performance WOSP '98**
Full text available: Additional Information: